

CLAIMS

What is claimed is:

1. A method for routing a packet ultimately intended for a destination made in an IP network by a receiving node, wherein all links between nodes have an assigned weight wherein the cost of a route is the sum of the weight of all links on that route, and wherein the shortest route path between the receiving node and the destination node is the route that has the lowest possible cost designated C , the method comprising:

determining whether the next link of the shortest route path from the receiving node to the destination node is congested; and

if the next link of the shortest route path from the receiving node to the destination node is congested:

identifying at least one node adjacent to the receiving node with a shortest route path between the adjacent node and the destination node having a cost less than C ;

determining if the link between the receiving node and each of the at least one identified adjacent nodes is congested; and

routing the packet to one of the at least one node if the link between the receiving node and the one of the at least one identified adjacent node is not congested.

2. The method for routing a packet of claim 1, wherein determining whether a link is congested comprises determining whether traffic on that link exceeds a predetermined fraction of the capacity of the link.

3. The method for routing a packet of claim 2, wherein the predetermined fraction of the capacity of the link comprises one-half.

4. The method for routing a packet of claim 3 further comprising, if the next link of the shortest route path from the receiving node to the destination node is congested:

determining which of the at least one identified adjacent node with a link between the identified adjacent node and the receiving node that is not congested has the lowest cost route to the destination node; and wherein:

routing the packet to one of the at least one identified adjacent node comprises routing the packet to the adjacent node determined to have the lowest cost route to the destination node.

5. A method for routing a packet ultimately intended for a destination node in an IP network by a receiving node, wherein all links between nodes have an assigned weight, wherein the cost of a route is the sum of the weight of all links on that route, wherein the minimum weight of an inter-PoP link is designated W_{\min} , wherein the maximum weight of an intra-PoP link is designated w_{\max} , wherein $W_{\min} > w_{\max}$, and wherein the shortest route path between the receiving node and the destination node is the route that has the lowest possible cost designated C , the method comprising:

determining whether the next link of the shortest route path from the receiving node to the destination node is congested; and

if the next link of the shortest route path from the receiving node to the destination node is congested:

identifying at least one node adjacent to the receiving node with a shortest route path between the adjacent node and the destination node not greater than $C + w_{\max}$ and with a cost between the at least one adjacent node and the receiving node less than w_{\max} ;

determining if the link between the receiving node and each of the at least one identified adjacent nodes is congested; and

routing the packet to one of the at least one identified adjacent node if the link between the receiving node and the one of the at least one identified adjacent node is not congested.

6. The method for routing a packet of claim 5, wherein determining whether a link is congested comprises determining whether traffic on that link exceeds a pre-determined fraction of the capacity of the link.

7. The method for routing a packet of claim 6, wherein the predetermined fraction of the capacity of the link comprises one-half.

8. The method for routing a packet of claim 7 further comprising, if the next link of the shortest route path from the receiving node to the destination node is congested:

determining which of the at least one identified adjacent node with a link between the identified adjacent node and the receiving node that is not congested has the lowest cost route to the destination node; and wherein:

routing the packet to one of the at least one identified adjacent node comprises routing the packet to the adjacent node determined to have the lowest cost route to the destination node.

9. A method for deflecting the routing of a packet around congestion in an IP network, the packet to be routed by a receiving node and ultimately intended for a destination node, wherein each link connecting a pair of nodes in the IP network is assigned a weight, wherein the cost of a route is the sum of the weights of all links in that route, wherein the shortest route path between a pair of nodes is the route between that pair of nodes with the lowest possible cost, wherein the smallest weight of an inter-PoP link is W_{\min} , wherein the largest weight of an intra-PoP link is w_{\max} , wherein $W_{\min} \gg w_{\max}$, wherein the cost of the shortest route path between the receiving node and the destination node is C , and wherein the next link of the shortest route path between the receiving node and the destination node is congested, the method comprising:

identifying nodes adjacent to the receiving node connected to the receiving node by non-congested links;

determining if one of the identified adjacent nodes has a shortest route path between that adjacent node and the destination node with a cost less than $C - w_{\max}$ and, if so, routing the packet to that node and, if not:

determining if one of the identified adjacent nodes:

has a cost no more than w_{\max} from the receiving node;

has a shortest route path between that adjacent node and the destination node with a cost no greater than $C + w_{\max}$; and

has a next link of the shortest route path between that adjacent node and the destination node that is not with the receiving node; and if so, routing the packet to that adjacent node.

10. The method for deflecting the routing of a packet of claim 9, wherein identifying nodes adjacent to the receiving node connected to the receiving node by non-congested links comprises determining that the link is non-congested if its traffic is below a predetermined fraction of the capacity of the link.

11. The method for deflecting the routing of a packet of claim 10, wherein the predetermined fraction of the capacity of the link is one half.

12. A method for deflecting the routing of a packet around congestion in an IP network, the IP network comprising a plurality of PoPs each with at least one router and a plurality of links connecting the routers, the packet to be routed by a receiving router and ultimately intended for a destination router, wherein each link connecting a pair of routers in the IP network is assigned a weight, wherein the cost of a route is the sum of the weights of all links in that route, wherein the shortest route path between a pair of routers is the route between that pair of routers with the lowest possible cost, wherein the smallest weight of a link connecting routers in different PoPs is W_{\min} , wherein the largest weight of a link connecting routers in the same PoP is w_{\max} , wherein $W_{\min} \gg w_{\max}$, wherein n denotes the number of routers in a PoP, wherein $W_{\min} > (n-1)w_{\max}$, and wherein the next link of the shortest route path between the receiving node and the destination node is congested, the method for deflecting the routing of the packet comprising:

identifying routers adjacent to the receiving router connected to the receiving router by non-congested links;

determining if one of the identified adjacent routers has a shortest route path between that adjacent router and the destination router with a cost less than $C - (n-1)w_{\max}$ and, if so, routing the packet to that router and, if not:

determining if one of the identified adjacent routers:

has a cost no more than w_{\max} from the receiving router;

has a shortest route path between the receiving router and the destination router with a cost no greater than $C + w_{\max}$; and

has a next link of the shortest route path between that adjacent router and the destination router that is not with the receiving router; and if so, routing the packet to that adjacent router.

13. The method for deflecting the routing of a packet of claim 12, wherein identifying routers adjacent to the receiving router connected to the receiving router by non-congested links comprises determining that the link is non-congested if its traffic is below a predetermined fraction of the capacity of the link.

14. The method for deflecting the routing of a packet of claim 13, wherein the predetermined fraction of the capacity of the link is one half.

15. A machine readable media containing machine readable code for causing a router to perform a method for routing a packet ultimately intended for a destination node in an IP network, wherein all links between nodes have an assigned weight, wherein the cost of a route is the sum of the weight of all links on that route, wherein the minimum weight of an inter-PoP link is designated W_{\min} , wherein the maximum weight of an intra-PoP link is designated w_{\max} , wherein $W_{\min} > w_{\max}$, and wherein the shortest route path between the receiving node and the destination node is the route that has the lowest possible cost designated C , the method comprising:

determining whether the next link of the shortest route path from the receiving node to the destination node is congested; and

if the next link of the shortest route path from the receiving node to the destination node is congested:

identifying at least one node adjacent to the receiving node with a shortest route path between the adjacent node and the destination node not greater than $C + w_{\max}$ and with a cost between the at least one adjacent node and the receiving node less than w_{\max} ;

determining if the link between the receiving node and each of the at least one identified adjacent nodes is congested; and

routing the packet to one of the at least one identified adjacent node if the link between the receiving node and the one of the at least one identified adjacent node is not congested.

16. The machine readable media of claim 15, wherein determining whether a link is congested comprises determining whether traffic on that link exceeds a pre-determined fraction of the capacity of the link.

17. The machine readable media of claim 16, wherein the predetermined fraction of the capacity of the link comprises one-half.

18. The machine readable media of claim 17, wherein the method further comprises, if the next link of the shortest route path from the receiving node to the destination node is congested:

determining which of the at least one identified adjacent node with a link between the identified adjacent node and the receiving node that is not congested has the lowest cost route to the destination node; and wherein:

routing the packet to one of the at least one identified adjacent node comprises routing the packet to the adjacent node determined to have the lowest cost route to the destination node.

19. A machine readable media containing machine readable code for causing a router to perform a method for deflecting the routing of a packet around congestion in an IP network, the IP network comprising a plurality of PoPs each with at least one router and a plurality of links connecting the routers, the packet to be routed by a receiving router and ultimately intended for a destination router, wherein each link connecting a pair of routers in the IP network is assigned a weight, wherein the cost of a route is the sum of the weights of all links in that route, wherein the shortest route path between a pair of routers is the route between that pair of routers with the lowest possible cost, wherein the smallest weight of a link connecting routers in different PoPs is W_{\min} , wherein the largest weight of a link connecting routers in the same PoP is w_{\max} , wherein $W_{\min} \gg w_{\max}$, wherein n denotes the number of routers in a PoP, wherein $W_{\min} > (n-1)w_{\max}$, and wherein the next link of the shortest route path between the receiving node and the destination node is congested, the method for deflecting the routing of the packet comprising:

identifying routers adjacent to the receiving router connected to the receiving router by non-congested links;

determining if one of the identified adjacent routers has a shortest route path between that adjacent router and the destination router with a cost less than $C - (n-1)w_{\max}$ and, if so, routing the packet to that router and, if not:

determining if one of the identified adjacent routers:

has a cost no more than w_{\max} from the receiving router;

has a shortest route path between the receiving router and the destination router with a cost no greater than $C + w_{\max}$; and

has a next link of the shortest route path between that adjacent router and the destination router that is not with the receiving router; and if so, routing the packet to that adjacent router.

20. The machine readable media of claim 19, wherein identifying routers adjacent to the receiving router connected to the receiving router by non-congested links comprises determining that the link is non-congested if its traffic is below a predetermined fraction of the capacity of the link.

21. The machine readable media of claim 20, wherein the pre-determined fraction of the capacity of the link is one-half.